arXiv:1603.06068v1 [cs.DB] 19 Mar 2016

Measuring the Accuracy of Linked Data Indices

Thomas Gottron

Institute for Web Science and Technologies
University of Koblenz-Landau
Universitätsstraße 1
56070 Koblenz
gottron@uni-koblenz.de

Abstract: Being based on Web technologies, Linked Data is distributed and decentralised in its nature. Hence, for the purpose of finding relevant Linked Data on the Web, search indices play an important role. Also for avoiding network communication overhead and latency, applications rely on indices or caches over Linked Data. These indices and caches are based on local copies of the original data and, thereby, introduce redundancy. Furthermore, as changes at the original Linked Data sources are not automatically propagated to the local copies, there is a risk of having inaccurate indices and caches due to outdated information. In this paper I discuss and compare methods for measuring the accuracy of indices. I will present different measures which have been used in related work and evaluate their advantages and disadvantages from a theoretic point of view as well as from a practical point of view by analysing their behaviour on real world data in an empirical experiment.

1 Introduction

The Linked Open Data (LOD) movement implements Tim Berner-Lee's vision of a *Web of Data*. The principles underlying Linked Data enable data providers to model, interlink and publish their data on the Web in a distributed and decentralised way. On the consumer's side the Web-oriented technological basis (i.e. using RDF and HTTP) permits to easily make use of the data and integrate it into various applications. These benefits and advantages of the Linked Data idea lead to, on the one hand, more and more data providers to contribute to the Web of Data and, on the other hand, more and more developers to spawn new applications making use of this data. However, the growth and conceptual nature of Linked Data poses some technological challenges. The distributed nature of Linked Data calls for search engines and indices to provide data catalogues of what kind of data is available and where it can be found. Furthermore, some applications use data caches to avoid communication overhead when accessing data on the Web.

Such indices and caches correspond to local views on the data. These local views might be inaccurate in the sense that they do not reflect the state of the data at the original data sources. There are two main reasons for indices constituting an inaccurate view on the data. The first reason is rooted in the decentralised and dynamic nature of Linked Data.

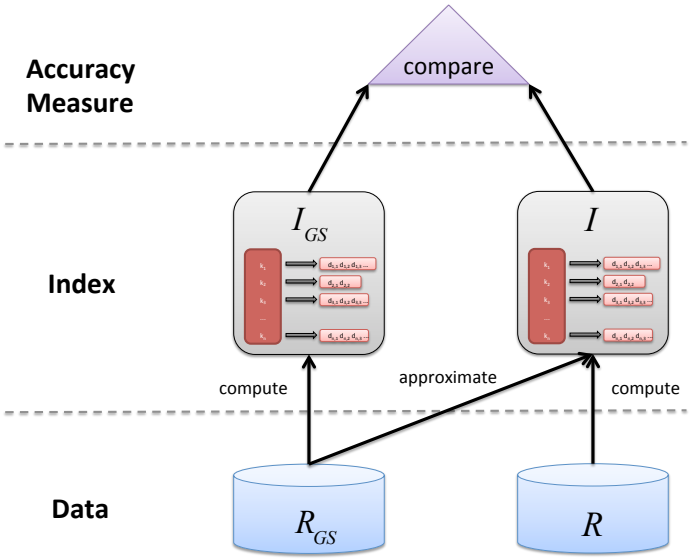

Figure 1: Inaccurate indices I can occur when computed over outdated data R or when computed in an approximative way over current data R_{GS}. A measure of accuracy needs to compare such an index I to a perfect, gold standard index I_{GS} computed over R_{GS} in a lossless way.

Changes at the original data sources are not automatically propagated to the applications and their indices and data caches. Accordingly a task which needs to be addressed in the context of Linked Data applications is the active maintenance of indices. This means, that applications have to synchronise their local view on the data with the data at the origin [UHH+10, DKG14]. A second reason are applications which build their indices only in an approximate way [KGSS12]. This happens, for instance, to achieve scalability and to be able to efficiently process large volumes of Linked Data. Figure 1 illustrates these two paths which can lead from a Linked Data set R_{GS} and an older version of this data set R to an index I which is inaccurate compared to a *perfect* index I_{GS}.

In both cases, there is a certain tradeoff between index accuracy and the commitment of limited resources. When updating and maintaining an index, the limited resource is typically network bandwidth. Spending more effort on updating an index requires more bandwidth but will yield indices of higher accuracy. Likewise, the approximate computation of indices can usually be influenced by a parameter which trades index accuracy for required computational resources. Hence, given the limitations of available resources the providers of LOD indices will need to strike a balance between the effort for building and maintaining an index and the quality of service, i.e. the accuracy of their index.

While measuring the costs for computation and network bandwidth has been addressed in many other fields, measures for the accuracy of a Linked Data index are less established. In related work, different methods for measuring index accuracy have been described and applied. But, rarely more than one approach is used. This makes it difficult to compare results. Thus, the idea of this paper is to compare different measures for index accuracy

on a theoretical and practical level. This should help in judging and comparing existing results as well as to support the choice of measures to use in future research work in this field. In more detail, this paper will make three contributions:

Survey of Measures: Following a review of related work, this paper provides an extensive overview of approaches and methods used to measure the accuracy of index structures over Linked Data. All methods are presented in a unified and formal way for ease of comparison.

Theoretical Comparison: Considering the general task of RDF indices and the specific Web based setting of Linked Data, this paper compares approaches and methods to measure the accuracy of LOD indices w.r.t their theoretical limitations and advantages.

Practical Comparison: By using an established data set of evolving Linked Data the paper analyses the approaches and methods to measure the accuracy of LOD indices w.r.t their behaviour in practice. In particular it presents an analysis of the correlation of the measures to understand how far similar or different notions of index accuracy are captured by the discussed measures.

As a next step we look at related work in Section 2 to get an overview of the context in which LOD indices are used and evaluated for their accuracy. This section will also provide a first and brief review of the methods used to measure index accuracy. Subsequently, Section 3 will present an abstract and formal representation for Linked Data indices. This formalisation allows to abstract from concrete implementations and serves as basis for the unified definition of index accuracy measures in Section 4. The theoretic analysis and comparison of the measures is presented in Section 5 and the practical and empirical comparison in Section 6. The findings are discussed in 7 and the paper is concluded with a summary and an outlook at future work.

2 Related Work

In recent years, various index models over LOD have been proposed. Many of them focus on specific aspects of the data or are dedicated to support application specific tasks. When looking at the RDF basis of LOD, one has to consider also the works on indices for RDF triple stores, such as Hexastore [WKB08] or RDF3X [NW10]. These indices are intended for optimising access to a single, centrally managed data storage solution. In this case, accuracy of the index is not an issue as all changes in the data are under control of the storage solution and are reflected in the index immediately.

More specific to LOD are indices for optimising federated queries [NW09], on-demand queries on the Web [HHK+10] or looking up data sources relevant to particular schema patterns [KGSS12]. However, most of these approaches do not deal with index accuracy either. Their focus is more on how to implement or make use of the index in specific scenarios.

So far, only few publications address the issue of Linked Data index accuracy. This only happens in a context where the dynamics of data is explicitly addressed [KAU+13, DSGG13, DGSG14], where implementations trade a loss of index accuracy for an efficient and scalable index computation [KGSS12, GP13] or where the accuracy of indices is under investigation itself [GG14, Got14, GP13]. These methods used will be reviewed and explained in more detail in Section 4.

In the context of the "classical" Web of documents, there is some work on index accuracy of search engines [LG99]. Such analysis is performed by comparing the coverage of relevant web documents for specific queries. The results of several individual search engines can be compared to the union of their results as well as to a separately obtained collection of relevant documents which serves as ground truth. Other works attempt to measure age and freshness of an index [CGM00]. Incorporating time information into a change prediction allows for more efficient and effective synchronisation plans between Web based data sources. Yet another direction of research addressed the question of guarantees of freshness of cached copies for web documents [BC00]. However, measuring the freshness and accuracy of search indices for Web documents is different from the accuracy of Linked Data indices insofar as the indices are all of the same type, namely mapping keywords to documents. In the context of Linked Data indices can be of different types and address specific information needs which are directly encoded in the index structure [GG14].

3 Abstract Index Models for Linked Data

As indicated in the previous section, there is a wide range of different index models over LOD. In this paper we will not look at specific implementations of indices, but rather operate on an abstract level. Therefore, we now briefly recall a formalisation of abstract index models over Linked Data [GG14]. This formalisation will serve as basis for a generic and unified definition of accuracy measures in Section 4.

On the LOD cloud we can assume data items to be in the form of NQuads [Car14]. In an NQuad (s, p, o, c) the entries s, p, and o correspond to the subject, predicate and object of the RDF triple statement. The entry c provides the context, i.e. the data source on the Web where this information has been published.

Thus, we define an index model for LOD over a data set R of (s, p, o, c) NQuads. Depending on the application scenario an index will typically not serve to store all information contained in the NQuads. Rather it will define a derived set D of managed data items which are of interest in the scenario and typically constitute a restriction of the quads to smaller tuples. Such restrictions can be, for instance, the RDF triples or even single entries, e.g. the subject or the context URIs. Furthermore, an index model has to define a set $\mathcal{K}$ of key elements which are used to lookup and retrieve data items. These key elements are used as domain for a selection function $\sigma : \mathcal{K} \to \mathcal{P}(D)$ to select a subset of the data items in the index.

Eventually, an abstract index model is defined as a tuple $(D, \mathcal{K}, \sigma)$ of the stored data items D, the key elements $\mathcal{K}$ used for the lookup index and the selection function σ to

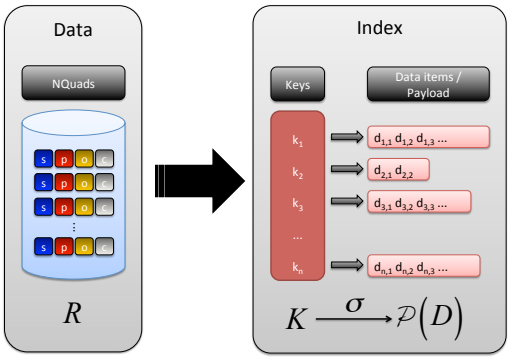

Figure 2: A data set R of NQuads and the elements of an index model over this data.

retrieve data from the index. Figure 2 illustrates the elements of an abstract index model I computed over a data set R.

4 Measuring Accuracy

The formal definition of abstract index models in Section 3 enables us to now formalise measures for index accuracy in a unified way. As already indicated in Figure 1, measuring the accuracy of an index will be based on a *perfect* index which is entirely accurate. In the following, this *gold standard* index will be referred to as $I_{\text{GS}} = (D_{\text{GS}}, \mathcal{K}_{\text{GS}}, \sigma_{\text{GS}})$ which is built over a data set R_{GS}. The potentially inaccurate index for which we want to determine its accuracy will be referred to as $I = (D, \mathcal{K}, \sigma)$. Note, that it is not necessary to distinguish between I being inaccurate because it was built over an outdated data set R or because it was computed in an approximative way over the gold standard data set R_{GS}.

The presented measures are divided into four families, based on their methods and underlying ideas for measuring accuracy: (1) index agnostic measures, (2) measures based on the overlap of key elements, (3) distribution based measures and (4) retrieval based measures.

4.1 Index Agnostic Measures

A direct way to measure the potential impact of data changes on indices is to simply compare the differences in the data itself. As indicated in Figure 3 such an accuracy measure would accordingly ignore the index and operate only on the original input data.

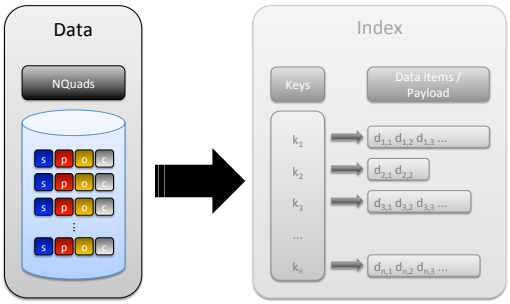

Figure 3: Index agnostic accuracy measures consider only the original data set.

One established metric for comparing the data sets is the *Jaccard similarity*. Applying it to R_{GS} and R, it is defined as:

$$\textit{Jaccard}(R_{\text{GS}}, R) = \frac{|R_{\text{GS}} \cap R|}{|R_{\text{GS}} \cup R|} \tag{1}$$

A higher similarity value indicates a larger overlap between the data sets while a low value indicates a stronger deviation, i.e. change in the data.

Such an index agnostic measures can always be computed over two versions of a data set at different points in time. Therefore, it is independent of the index model applied. While this might be suitable to measure the evolution of the data in general, it does not make any statement about the impact on specific index models used in specific use case scenarios. Furthermore, it requires the availability of the full raw data sets and is by design not applicable in settings were indices are computed in an approximative way over the gold standard data set R_{GS}.

4.2 Overlap of Key Elements

A relatively simple approach for comparing indices themselves is to look at the key elements (cf. Figure 4). These are the elements used to retrieve information from the index and, thus, provide the primary access point to the contained data. Hence, the degree of how far an index reflects the right key elements can already give some insights into its accuracy. Note, that this type of measure will not consider at all the data elements. In particular, it might assert an index a perfect accuracy even if the index provides wrong results.

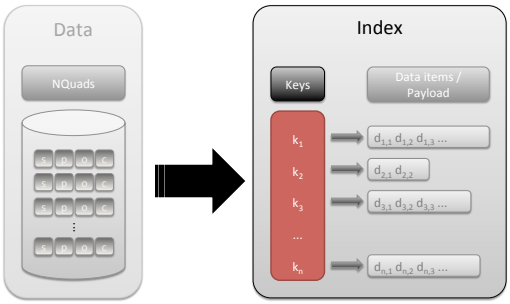

Figure 4: Accuracy measures based on the key elements do not make use of any information about the data items.

Technically, this approach comes down to comparing two sets: $\mathcal{K}_{\mathrm{GS}}$ and $\mathcal{K}$. There are various measures for comparing sets. In the context of measuring the accuracy of LOD indices, also here the *Jaccard similarity* has been used [GG14]:

$$Jaccard(\mathcal{K}_{\mathrm{GS}}, \mathcal{K}) = \frac{|\mathcal{K}_{\mathrm{GS}} \cap \mathcal{K}|}{|\mathcal{K}_{\mathrm{GS}} \cup \mathcal{K}|} \tag{2}$$

Again, a higher similarity value indicates a larger overlap between the sets of key elements while a low value indicates a stronger deviation. In this way we can get an impression of how stable is the set of elements used for indexing in the different indexing approaches.

A further measures to operate on the set of key elements is the asymmetric *recall*, which measures which fraction of the gold standard key set is covered in an inaccurate index.

$$r(\mathcal{K}_{\mathrm{GS}}, \mathcal{K}) = \frac{|\mathcal{K}_{\mathrm{GS}} \cap \mathcal{K}|}{|\mathcal{K}_{\mathrm{GS}}|} \tag{3}$$

The rational for using this variation is that key elements in an inaccurate index which are not available in the gold standard might not be used to retrieve information as well. However, this rational is debatable as a query with one of this key elements missing in the gold standard will lead to retrieving false positive information from the inaccurate index, a fact which is neglected by recall.

4.3 Distribution Based Measures

One option to include some more information in the evaluation of index accuracy is pursued by distribution based approaches [GG14, Got14]. The advantage is that these approaches take into consideration the volume of data items retrieved for a given key element. This idea is visualised in Figure 5 by omitting the concrete data items assigned to a

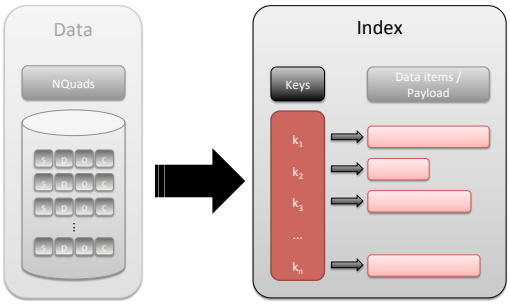

Figure 5: Distribution based measures compare distributions of the data over key elements.

key element. The volume of the data is used to model probabilistic distributions for which there are several well established and well understood metrics for comparison. However, they do not distinguish the data items but merely use their count information. Hence, also here there is a certain risk of asserting a perfect index accuracy while the index actually is not accurate.

Estimating a distribution over an index is based on determining how probable it is for an element to belong to one specific index key k and—conversely—the amount of data obtained when querying the index for this key element k. If we consider the distribution over an index $I = (D, \mathcal{K}, \sigma)$, this effectively corresponds to modelling a random variable X taking values of the key elements $\mathcal{K}$. The estimated density gives the distribution of this random variable X. This means we determine the probability $P(X = k)$ for each entry $k \in \mathcal{K}$ to be associated with a data item. To estimate the densities we can use the count information of data elements associated with the key elements in an index. This corresponds to using a maximum likelihood estimation to derive the probability, i.e.

$$P(X = k) = \frac{|\sigma(k)|}{\sum_{k' \in \mathcal{K}} |\sigma(k')|} \tag{4}$$

where $\sigma(k)$ indicates the result set obtained from an index when querying for a specific key element k.

As we consider inaccurate or outdated indices it is likely that the set of key elements does not match (see also Section 4.2). Thus, it can happen that certain key elements might be available in the perfect gold standard index, but not in an outdated or approximated index. When comparing densities over indices we need to consider the effect of such zero-size entries. Using a maximum likelihood estimation for the densities would lead to zero probabilities for certain events which renders comparison of densities impractical. We apply smoothing to overcoming zero probabilities. We make use of Lidstone smoothing which adds a small constant value of λ to all counts obtained for the number of results

$|\sigma(k)|$. The parameter λ is set to 0.5 in our experiments which has shown to provide good results in prior work [Got14].

The main idea of measuring index accuracy on the basis of data distributions is to compare their density function. Common metrics to compare density functions are cross entropy, Kullback-Leibler divergence and perplexity. Let us briefly review the definitions of these metrics and explain their interpretation.

Assuming two probability distributions $P(X)$ and $P_{\text{GS}}(X)$ for the inaccurate index I and the gold standard index I_{GS}. Then *cross entropy* $H(P_{\text{GS}}, P)$ is defined as:

$$H(P_{\text{GS}}, P) = -\sum_{k \in \mathcal{K} \cup \mathcal{K}_{\text{GS}}} P_{\text{GS}}(X = k) \log(P(X = k)) \tag{5}$$

In the context of compression theory, cross entropy can be interpreted as the average number of bits needed to encode events following the distribution P_{GS} based on an optimal encoding scheme derived from P. If the two distributions are equivalent, then cross entropy corresponds to the normal entropy $H(P_{\text{GS}})$. The entropy of P_{GS} also provides a lower bound for cross entropy. Based on this interpretation, the *Kullback-Leibler divergence* gives the deviation in entropy (or overhead in encoding) relative to the entropy for P_{GS} and is defined as:

$$D_{KL}(P_{\text{GS}}, P) = H(P_{\text{GS}}, P) - H(P_{\text{GS}}) \tag{6}$$

Therefore, if two distributions are equivalent, they have a Kullback-Leibler divergence of zero.

Perplexity, instead, provides an evaluation of a distribution by giving the number of events (in our case key elements) which under a uniform distribution would yield the same entropy value. As such it is considered to be more easily interpretable by humans than the somewhat abstract entropy values. Perplexity itself is defined over entropy values, though. Here we formulate it directly on the basis of cross entropy:

$$PP(P_{\text{GS}}, P) = 2^{H(P_{\text{GS}}, P)} \tag{7}$$

Perplexity is a standard metric for evaluating probabilistic models. The lower the perplexity is, the better a model explains observed data and the more truthful are its estimates of the probabilities. Furthermore, the interpretation of perplexity relative to the event space of key elements allows for a normalisation. The normalised perplexity PP_{norm} is defined as:

$$PP_{norm}(P_{\text{GS}}, P) = \frac{PP(P_{\text{GS}}, P)}{|\mathcal{K}_{\text{GS}}|} \tag{8}$$

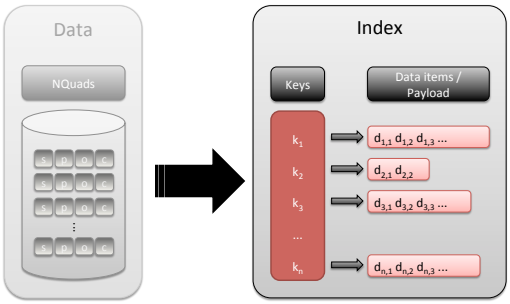

Figure 6: Retrieval based measures include the entire index in the evaluation.

4.4 Retrieval Based Measures

The third type of measures are directly evaluating the results of queries posed towards an index. *Precision* and *recall* are two such measures which are typically employed in Information Retrieval. While in general, this approach can be based on a fixed and predefined set of queries, such a set of queries might not be available in most cases. Furthermore, a fixed set of queries will always address only a certain fraction of the index. Thus, an alternative approach which has been followed in related work is to pose virtually all possible queries which could be asked over the gold standard as well as the inaccurate index. In this way, the measures make use of all information stored in the index as indicated in Figure 6.

For one specific query, i.e. a key item k, precision and recall are defined as follows:

$$p(k) = \frac{|\sigma(k) \cap \sigma_{\text{GS}}(k)|}{|\sigma(k)|} \tag{9}$$

$$r(k) = \frac{|\sigma(k) \cap \sigma_{\text{GS}}(k)|}{|\sigma_{\text{GS}}(k)|} \tag{10}$$

There is, however a distinction to be made with respect to the way of aggregating results over the set of all queries. Namely, there are micro- and macro-averages which can be computed.

Macro-average builds an average over all precision or recall values for all of the queries. Formally this gives (here for the example of precision):

$$p_{macro} = \frac{1}{|\mathcal{K} \cup \mathcal{K}_{\text{GS}}|} \sum_{k \in \mathcal{K} \cup \mathcal{K}_{\text{GS}}} \frac{|\sigma(k) \cap \sigma_{\text{GS}}(k)|}{|\sigma(k)|} \tag{11}$$

Micro-average, instead, aggregates the set sizes of the observed result sets and their intersection. Precision and recall are then computed on these aggregated count information. Formally, it is defined as:

$$p_{micro} = \frac{\sum_{k \in \mathcal{K} \cup \mathcal{K}_{\mathrm{GS}}} |\sigma(k) \cap \sigma_{\mathrm{GS}}(k)|}{\sum_{k \in \mathcal{K} \cup \mathcal{K}_{\mathrm{GS}}} |\sigma(k)|} \tag{12}$$

The difference of the two ways for computing averages can be seen as follows. Macro-average gives equal weight to each query. Hence a lot of good or bad performing queries which provide small result sets will strongly affect the overall score. Computing a micro-average, instead, considers also the size of the results set in the overall aggregated score. Larger result sets and their performance have a stronger impact on the overall result.

5 Theoretic Analysis of Measures

Let us consider the measures introduced above in Section 4 and compare them on a theoretic level. To this end we will consider the following four criteria and argue for and against the individual approches, independent of a concrete scenario, index model or data set.

Sensitivity to Data Changes Not all measures perfectly reflect the accuracy of an index. In some cases a change in the data and a loss of accuracy of an outdated index will not be captured by an accuracy measure. The higher the sensitivity of a measure the more precise and reliable it is in evaluating an index's ability to provide truthful results. In the context of this paper the theoretic sensitivity is judged in a relative and categorial way, by assigning the measures a low (–), medium (o) or perfect (+) sensitivity. A perfect sensitivity implies that each error in an inaccurate index will be detected.

Data Volume When computing the measures, one question is how much data needs to be obtained from a perfect index at the original data source. A smaller volume of data which needs to be transfered from separate indices makes it easier to compute the measures remotely and only request the needed information from the actual computing nodes running the index. Furthermore, the computational complexity of the presented measures is linear in the amount of data items to consider. Thus, the data volume indicates how long the computation of the measures will take. For this reason, the data volume will be given in big O notation based on the elements of abstract index models.

Normalised Value Range A normalised value range is favourable in an index accuracy measure. If a measure provides a normalised value range it is easier to compare its values across experiments. This can affect both: the comparison of accuracy of the same index over different data sets as well as the comparison of different indices

over the same data set. Having a normalised value range is a binary feature of an accuracy measure, i.e. a measure can either have a normalised range (√) or not (×).

Index Agnostic If a measure is index agnostic it can be computed independently of a concrete index model or implementation. The advantage is, that the computation of one single value over the data is sufficient to judge a (potential) impact on different types of indices. A conceptual disadvantage is, that index agnostic measures cannot be used to evaluate approaches performing an approximate index computation. Also being index agnostic is a binary feature, i.e. the measure is index agnostic (√) or not (×).

Table 1 summarises the theoretic evaluation of the measures from Section 4 w.r.t the above mentioned criteria. Within each family of measures, the characteristics are mostly the same. Typically, the measures have the same sensitivity, require the same volume of data and are all either index specific or index agnostic.

Judging from this theoretic point of view, the retrieval based measures seem favourable. After all they directly address the purpose of an index and measure the accuracy using the core functionality. Furthermore, they provide normalised values which renders them suitable for comparisons across data sets and types of indices. However, they require a large volume of data. Basically, all information of the entire index has to be considered for a comparison, which—depending on the type of index—might be as costly as comparing the entire underlying RDF data set. Under this aspect, the measures using the overlap of key elements and the distribution based measures need less information. They merely operate on the set of key elements and—in the case of the distribution based measures—a size estimation of how many elements are associated with a key element. Such size estimations can be realised using count operators over the existing indices, causing only a slight computational overhead. Their drawback is that they might not always capture perfectly the accuracy of an index.

6 Empirical Comparison of Measures

Following the theoretical evaluation of the measures we now consider how they perform in practice. The question we want to answer is which measures are closely correlated to each other in practice. The aim is to find out which measures seem to be redundant and do not need to be compared. This might help in reducing overhead in practical evaluation and in identifying efficient measures to approximate or even substitute more complex measures. In particular, a correlation analysis can also address the question how large is the effect of data changes in practice on the less sensitive key element and distribution based measures, i.e. how far they still agree with the measures having a perfect sensitivity.

To perform such an evaluation we need a set of indices (for the index specific measures) and a data set which undergoes realistic changes. As data set we can employ the weekly snapshots of a well defined part of the Linked Data cloud provided by the Dynamic Linked Data Observatory (DyLDO) [KAU+13]. The range of index models is wide and it is

Table 1: Theoretic analysis of measures for index accuracy

Measure	Sensitivity	Normalised range	Volume	Index agnostic
Index agnostic measures				
Jaccard (RDF triples)	+	√	$O(\|R\| + \|R_{\text{GS}}\|)$	√
Overlap of key elements				
Jaccard (Key elements)	–	√	$O(\|\mathcal{K}\| + \|\mathcal{K}_{\text{GS}}\|)$	×
Recall (Key elements)	–	√	$O(\|\mathcal{K}\| + \|\mathcal{K}_{\text{GS}}\|)$	×
Distribution based measures				
Cross-Entropy	o	×	$O(\|\mathcal{K}\| + \|\mathcal{K}_{\text{GS}}\|)$	×
KL-Divergence	o	×	$O(\|\mathcal{K}\| + \|\mathcal{K}_{\text{GS}}\|)$	×
Perplexity	o	×	$O(\|\mathcal{K}\| + \|\mathcal{K}_{\text{GS}}\|)$	×
Normalised Perplexity	o	√	$O(\|\mathcal{K}\| + \|\mathcal{K}_{\text{GS}}\|)$	×
Retrieval based measures				
Recall (macro-avg)	+	√	$O(\|\mathcal{D}\| + \|\mathcal{D}_{\text{GS}}\|)$	×
Precision (macro-avg)	+	√	$O(\|\mathcal{D}\| + \|\mathcal{D}_{\text{GS}}\|)$	×
Recall (micro-avg)	+	√	$O(\|\mathcal{D}\| + \|\mathcal{D}_{\text{GS}}\|)$	×
Precision (micro-avg)	+	√	$O(\|\mathcal{D}\| + \|\mathcal{D}_{\text{GS}}\|)$	×

beyond the scope of this paper to evaluate all types of index models. Thus, let us chose a few selected index models which cover a good range of index granularities and scope.

The indices are taken vom related work and operate on different types of data they return: (1) full triple statements (D_{triple}), (2) single URIs (D_{URI}) or (3) solely the source where relevant information can be found on the Linked Data cloud ($D_{context}$). In the following, the index models are described briefly along with a formal definition for clarity. More details can be found in original publications.

Subject Index: This type of index simply uses URIs as key elements and its selection function returns all triple statement containing a given URI in the subject position. Using the abstract notation introduced in Section 3, such an index can be formalised as $I_S := (D_{triple}, \mathcal{K}_S, \sigma_S)$ where:

- Key elements: $\mathcal{K}_S := \{s \in U \mid \exists p, o, c : (s, p, o, c) \in R\}$
- Selection function: $\sigma_S(k) := \{(s, p, o) \mid s = k\}$

RDF Type Index: Making use of the specific semantics behind rdf:type statements, this index is used to look up all entities (i.e. URIs) which are of a particular RDF class type [GG14]. Formally it can be defined as $I_T := (D_{URI}, \mathcal{K}_T, \sigma_T)$, with:

- Key elements: $\mathcal{K}_T := \{o \mid \exists s, c : (s, \text{rdf:type}, o, c) \in R\} \cup \{s \mid \exists c : (s, \text{rdf:type}, \text{rdfs:Class}, c) \in R\}$
- Selection function: $\sigma_T(k) := \{s \mid \exists c : (s, \text{rdf:type}, k, c) \in R\}$

RDF Type Set (TS) Index: Extending the RDF type class index, this index gives all entities for a specific set of class types, i.e. all entities which satisfy being of exactly all types given in the set (and of no other type) [KGSS12]. The formal definition is given by $I_{TS} := (D_{URI}, \mathcal{K}_{TS}, \sigma_{TS})$, where:

- Key elements: $\mathcal{K}_{TS} := \mathcal{P}(\mathcal{K}_T)$
- Selection function: $\sigma_{TS}(k) := \{s \mid (\forall t \in k : (\exists c : (s, \text{rdf:type}, t, c) \in R)) \wedge \forall (s, \text{rdf:type}, o, c) \in R : (o \in k)\}$

Property Set (PS) Index: Using sets of predicates—also known as characteristics sets—this index retrieves entities based on the properties with which they have been described in RDF [NW09]. Formally it can be defined as $I_{PS} := (D_{URI}, \mathcal{K}_{PS}, \sigma_{PS})$, with:

- Key elements: $\mathcal{K}_{PS} := \mathcal{P}(\mathcal{K}_p)$
- Selection function: $\sigma_{PS}(k) := \{s \mid (\forall p \in k : (\exists o, c : (s, p, o, c) \in R)) \wedge \forall (s, p, o, c) \in R : (p \in k)\}$

Extended Characteristic Set (ECS) Index: This index essentially combines the TS and PS index and characterises Linked Data entities by both: their types and properties [DSGG13]. Formally it is defined by $I_{ECS} := (D_{URI}, \mathcal{K}_{ECS}, \sigma_{ECS})$, where:

- Key elements: $\mathcal{K}_{ECS} := \mathcal{P}(\mathcal{K}_P \cup \mathcal{K}_T)$
- Selection function: $\sigma_{ECS}(k) := \{s \mid (\forall p \in k \cap \mathcal{K}_P : (s \in \sigma_{PS}(p))) \wedge (\forall t \in k \cap \mathcal{K}_T : (s \in \sigma_{TS}(t)))\}$

SchemEX Index over data sources: The SchemEX index combines induced schematic information about the types of subject URIs, their properties and the types of object URIs they are linked to [KGSS12]. To this end it re-uses the concepts of type clusters for the domain and range of observed property sets. Those concepts are combined using a restricted bi-simulation over property sets and a stratification of the entities in equivalence classes of type sets. Its formal definition is given by $I_{SchemEX} := (D_{context}, \mathcal{K}_{SchemEX}, \sigma_{SchemEX})$, where:

- Key elements: $\mathcal{K}_{SchemEX} := \mathcal{P}(\mathcal{K}_{TS} \times \mathcal{P}(\mathcal{K}_{PS} \times \mathcal{K}_{TS}))$
- Selection function: $\sigma_{SchemEX}(k = (ts, E)) := \{c \mid \exists s \in \sigma_{TS}(ts) \wedge \forall (ps, ts_2) \in E : (s \in \sigma_{PS}(ps) \wedge \exists o \in \sigma_{TS}(ts_2) : (\forall p \in ps : ((s, p, o, c) \in R))\}$

For the empirical evaluation we computed these indices over 77 snapshots provided by the DyLDO data set. Assuming an index has been computed over the initial snapshot, we evaluate it accuracy compared to the later on versions of the data sets. This corresponds to fixing an index over the initial data set and comparing it to gold standard indices over later data sets. All measures have been computed for all of the evaluated index models. This gives us for each update of the data set one observation of the accuracy measures. The results of a spearman rank correlation analysis between all pairs of measures over this data is visualised in Figure 7.

There are several interesting observations to be made. First of all we observe two index models where the correlations are much stronger in general: the subject index and SchemEX. All other analysed index models exhibit lower correlation values. The explanation for this behaviour is the granularity of the indices. Both, the subject index and

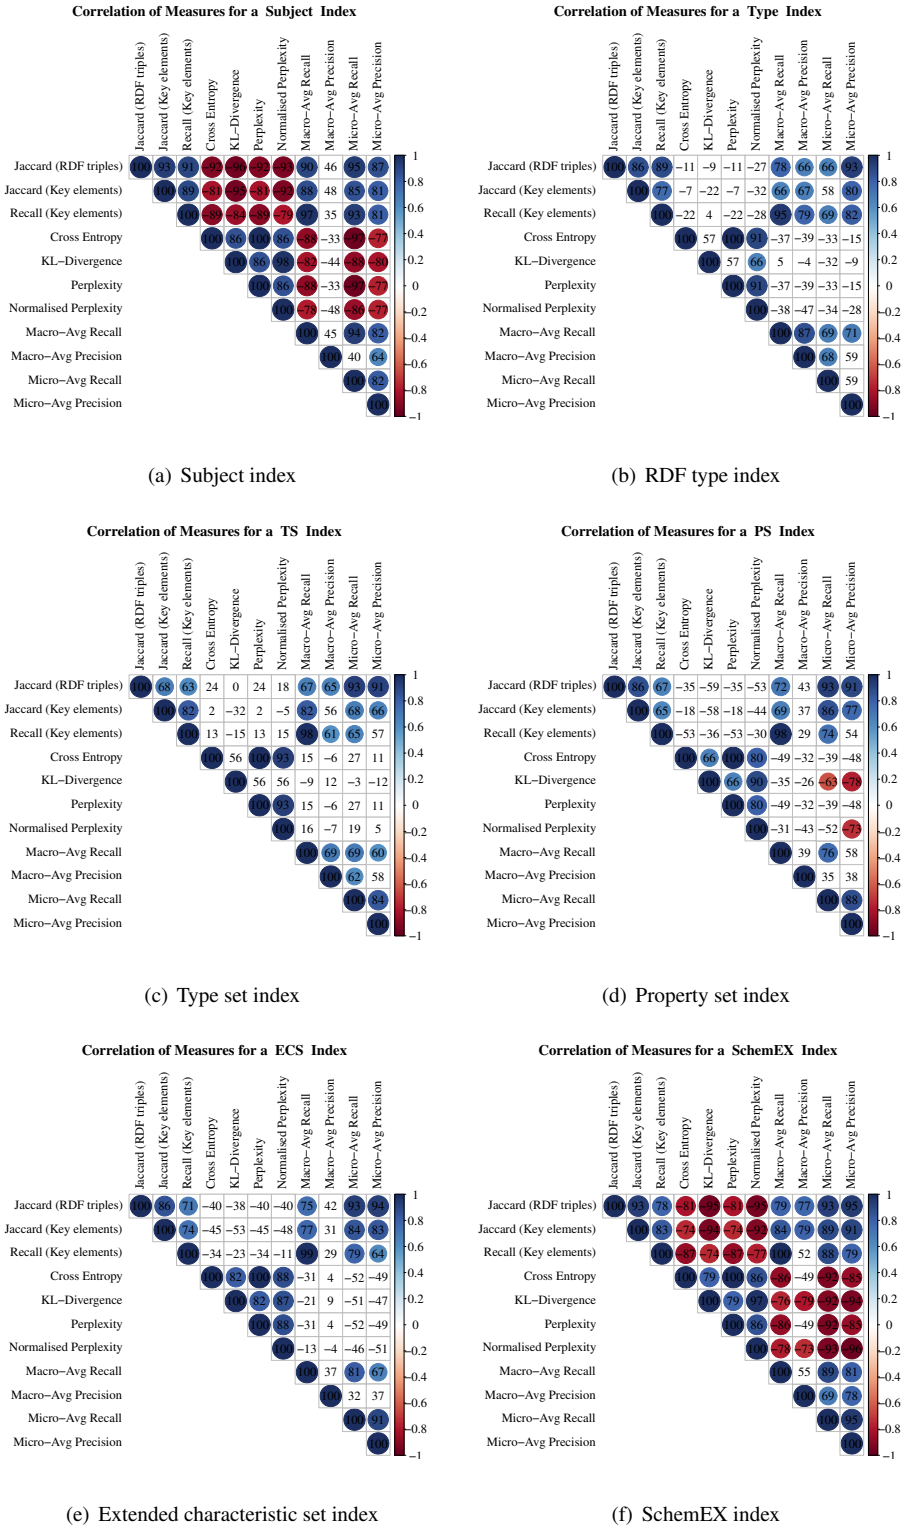

(a) Subject index

(b) RDF type index

(c) Type set index

(d) Property set index

(e) Extended characteristic set index

(f) SchemEX index

Figure 7: Correlation of accuracy measures on different indices over an evolving data set.

SchemEX, generate a relatively large number of key elements over the analysed data set. Furthermore, they have many key elements in the index which provide only one element entry upon request. This means on the one hand, that measuring the accuracy of a specific query is often reduced to a binary task: either the index provides the correct element or not. Therefore, the retrieval based and key element based measures are strongly correlated in a positive way. On the other hand, as the change of a data element being assigned to a different key element is typically caused by a single triple, also the Jaccard index over triples correlates strongly with the measures. Also the changes in the distribution are affected by this. Changes for key elements with one entry typically leads to a strong shift in the overall distribution, reducing a MLE probability to zero[1]. This shift in probability mass is recognised by the distribution based measures. The observation that the correlation here is highly negative comes from the fact that for those measures a high value indicates a strong change, while the other measures use high values to indicate no or only weak changes.

For all other analysed index models, we have a different overall behaviour. Typically, the measures within each family are correlated (though to a lower extent than for the subject index or SchemEX). Furthermore, the key element based measures correlate strongly with the retrieval based measures and the Jaccard index over the sets of RDF triples. The distribution based measures, instead seem to measure something else. Beyond a few exception, no strong correlation can be observed here.

Looking at individual measures we can also observe some interesting behaviour: Across all types of index models macro-average precision seems to be not or only weakly correlated to most other measures. Hence, the different way of aggregating precision seems to have a strong impact. The micro-average show a better sensitivity of the accuracy of queries with large results sets. This seems to make a distinction with other evaluation measures over most analysed index models. The perfect Spearman correlation between Cross-Entropy and Perplexity, instead, is conceptual as Perplexity is only a rank-preserving non-linear scaling of Cross-Entropy.

7 Discussion

Given the analysis in the previous two section, we can observe that a look at the key elements might already give a good indication for the accuracy of an index. While from a theoretic point of view they might not be very sensitive to changes they correlated well with other, more sensitive measures in the empiric evaluation. In particular they reflect quite well the retrieval based accuracy measures. This observations holds for all analysed index models.

Additionally, in many cases the distribution based measures provide further insights. Given that the volume of data needed to compute these measures is of the same order as for the key elements, such an analysis causes little overhead.

[1] Effectively the probability has a small value, slightly higher than zero due to the use of Lidstone smoothing as explained in Section 4.3.

In general, the empirical observations in the previous section were made only over a single data set. However, given that this data set has been carefully designed to give a representative excerpt of the LOD cloud, we might want to conjecture that density based measures as well as the measures based on key elements suffice to judge the quality of Linked Data indices. As information on the distributions over key elements in Linked Data index structures is sufficient to compute both types of measures, they are also attractive from a computational point of view.

8 Summary and Conclusions

In this paper I investigated the question of how to measure the accuracy of index structures and data caches which are built in an approximative way or over evolving Linked Data sets. I gave an overview of different approaches for such measures, compared them on a theoretical and empirical level. One observation is that information about the data distribution over key elements seems to provide good insights into index accuracy without the need to access all data in an index.

This observation also motivates a roadmap for future work. Operating on samples of Linked Data can provide good estimates of these distributions [Got14]. Accordingly an interesting question is how suitable are sampling based strategies for evaluating index accuracy. This means, that instead of asking all possible queries, we consider only a random sample of queries to obtain distributions. This might speed up computation or can at least provide certain confidence intervals about what the true accuracy of an index might be.

References

[BC00] Brian E Brewington and George Cybenko. How dynamic is the Web? *Computer Networks*, 33(1):257–276, 2000.

[Car14] Gavin Carothers. RDF 1.1 N-Quads. W3C Recommendation, Feb. 2014. (accessed 14 March 2014).

[CGM00] Junghoo Cho and Hector Garcia-Molina. Synchronizing a Database to Improve Freshness. In *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data*, SIGMOD '00, pages 117–128, New York, NY, USA, 2000. ACM.

[DGSG14] Renata Dividino, Thomas Gottron, Ansgar Scherp, and Gerd Gröner. From Changes to Dynamics: Dynamics Analysis of Linked Open Data Sources. In *PROFILES'14: Proceedings of the Workshop on Dataset ProfiIling and Federated Search for Linked Data*, 2014. (to appear).

[DKG14] Renata Dividino, Andre Kramer, and Thomas Gottron. An Investigation of HTTP Header Information for Detecting Changes of Linked Open Data S ources. In *ESWC'14: Proceedings of the Extended Semantic Web Conference*, 2014. (to appear).

[DSGG13] Renata Dividino, Ansgar Scherp, Gerd Gröner, and Thomas Gottron. Change-a-LOD: Does the Schema on the Linked Data Cloud Change or Not? In *COLD'13: International Workshop on Consuming Linked Data*, 2013.

[GG14] Thomas Gottron and Christian Gottron. Perplexity of Index Models over Evolving Linked Data. In *ESWC'14: Proceedings of the Extended Semantic Web Conference*, pages 161–175, 2014.

[Got14] Thomas Gottron. Of Sampling and Smoothing: Approximating Distributions over Linked Open Data. In *PROFILES'14: Proceedings of the Workshop on Dataset Profiling and Federated Search for Linked Data*, 2014. (to appear).

[GP13] Thomas Gottron and Rene Pickhardt. A Detailed Analysis of the Quality of Stream-Based Schema Construction on Linked Open Data. In Juanzi Li, Guilin Qi, Dongyan Zhao, Wolfgang Nejdl, and Hai-Tao Zheng, editors, *Semantic Web and Web Science*, Springer Proceedings in Complexity, pages 89–102. Springer New York, 2013.

[HHK+10] Andreas Harth, Katja Hose, Marcel Karnstedt, Axel Polleres, Kai-Uwe Sattler, and Jürgen Umbrich. Data summaries for on-demand queries over linked data. In *Int. Conf. on World wide web*, pages 411–420. ACM, 2010.

[KAU+13] Tobias Käfer, Ahmed Abdelrahman, Jürgen Umbrich, Patrick O'Byrne, and Aidan Hogan. Observing Linked Data Dynamics. In Philipp Cimiano, Oscar Corcho, Valentina Presutti, Laura Hollink, and Sebastian Rudolph, editors, *The Semantic Web: Semantics and Big Data*, volume 7882 of *Lecture Notes in Computer Science*, pages 213–227. Springer Berlin Heidelberg, 2013.

[KGSS12] Mathias Konrath, Thomas Gottron, Steffen Staab, and Ansgar Scherp. SchemEX—Efficient Construction of a Data Catalogue by Stream-based Indexing of Linked Data. *Web Semantics: Science, Services and Agents on the World Wide Web*, 16(5):52 – 58, 2012. The Semantic Web Challenge 2011.

[LG99] Steve Lawrence and C Lee Giles. Accessibility of information on the web. *Nature*, 400(6740):107–107, 1999.

[NW09] Thomas Neumann and Gerhard Weikum. Scalable join processing on very large RDF graphs. In *Int. Conf. on Management of data*, pages 627–640. ACM, 2009.

[NW10] Thomas Neumann and Gerhard Weikum. The RDF-3X engine for scalable management of RDF data. *The VLDB Journal*, 19(1):91–113, 2010.

[UHH+10] Jürgen Umbrich, Michael Hausenblas, Aidan Hogan, Axel Polleres, and Stefan Decker. Towards Dataset Dynamics: Change Frequency of Linked Open Data Sources. In *LDOW*, 2010.

[WKB08] Cathrin Weiss, Panagiotis Karras, and Abraham Bernstein. Hexastore: sextuple indexing for semantic web data management. *Proc. VLDB Endow.*, 1(1):1008–1019, August 2008.